Rochester Center for
Economic Research

Individual and Aggregate Labor Supply in a Heterogeneous Agent Economy with Intensive and Extensive Margins

Yongsung Chang Sun-Bin Kim Kyooho Kwon Richard Rogerson

Working Paper No. 583
July 2014

University of
Rochester

Individual and Aggregate Labor Supply in a Heterogeneous Agent Economy with Intensive and Extensive Margins*

Yongsung Chang
University of Rochester
Yonsei University

Sun-Bin Kim
Yonsei University

Kyooho Kwon
Korea Development Institute

Richard Rogerson
Princeton University
NBER

July 1, 2014

Abstract

We develop a heterogeneous-agent general equilibrium model that incorporates both intensive and extensive margins of labor supply. A nonconvexity in the mapping between time devoted to work and labor services distinguishes between extensive and intensive margins. We consider calibrated versions of this model that differ in the value of a key preference parameter for labor supply and the extent of heterogeneity. The model is able to capture the salient features of the empirical distribution of hours worked, including how individuals transit within this distribution. We then study how the various specifications influence labor supply responses to aggregate technology shocks. We find that abstracting from intensive margin adjustment has large effects on the volatility of aggregate hours even if intensive margin fluctuates relatively little.

Keywords: Hours, Employment, Cross-section, Business Cycles

*Correspondence: Chang: Department of Economics, University of Rochester and Yonsei University, yongsung.chang@gmail.com; Kim: Department of Economics, Yonsei University, sunbin.kim@yonsei.ac.kr; Kwon: Korea Development Institute, kwonkh98@gmail.com; Rogerson: Department of Economics, Princeton University, rdr@princeton.edu.

We thank Richard Blundell for comments as well as conference participants at Yonsei University, Institute for Fiscal Studies and the 2012 SED. Chang and Rogerson gratefully acknowledge the support from the WCU (R33-10005) Program through the National Research Foundation of Korea. Rogerson also thanks the NSF for financial support. Kim's research was supported by the National Research Foundation of Korea Grant funded by the Korean Government (NRF-2011-330-B00055).

1. Introduction

Early representative household models used to study business cycle fluctuations (e.g., Kydland and Prescott (1982)) assumed a household that smoothly adjusts labor supply. An immediate controversy revolved around the fact that these models were calibrated so as to give the representative household a Frisch labor supply elasticity larger than those estimated on micro data by researchers such as MaCurdy (1981) and Altonji (1986). Heckman (1984) argued that this controversy was somewhat misdirected given the extensive margin was the dominant margin of adjustment during US business cycles and both the business cycle models and the micro estimation exercises abstracted from this margin. It is now understood that the labor supply elasticity of a representative household should capture adjustment along both the intensive and extensive margins. However, starting with the study of Hansen (1985), it has become common for macroeconomic analyses to assume that all labor supply adjustment occurs along the extensive margin. The usual motivation for this assumption is that between two thirds and three quarters of business cycle fluctuations in aggregate hours are due to changes in employment rather than hours per worker. A natural question, but one that has not been asked in the literature, is whether abstracting from the intensive margin is a (relatively) harmless simplification for understanding cyclical fluctuations in the labor market. More generally, how does the presence of an intensive margin affect the implied aggregate labor supply elasticity. In this paper we demonstrate that including an empirically reasonable channel of intensive margin choice has important quantitative implications for labor market fluctuations. In particular, we show that even if fluctuations along the intensive margin are very small, their presence can have large aggregate effects on overall fluctuations in the labor market. Relative to a model that abstracts from the intensive margin, we find that a model with an active intensive margin implies a much more compressed range for aggregate elasticities when considering empirically reasonable variation in both micro Frisch elasticities and the extent of heterogeneity.

The first step in our analysis is to develop a steady state model that can match the salient features of individual labor supply along the intensive and extensive margins. To do this we embed the nonconvex production model of Prescott et al. (2009) into the indivisible labor framework of Chang and Kim (2007) that features idiosyncratic shocks and incomplete markets. We assess the ability of this model to account for various steady state observations, including the distribution of hours of work across individuals, the transition of individuals in the hours of work distribution over time, and the distribution of labor earnings and wealth. Although our model is parsimonious, it is able to account for many stylized facts in the data. We note two novel aspects to our steady state calibration exercise. First, we argue that the cross-sectional distribution of hours of work can serve as useful information regarding the empirically relevant amount of heterogeneity. Interestingly, based on this measure, we need a degree of heterogeneity that is roughly double the amount captured by estimates of idiosyncratic wage shocks. It turns out that the extent of heterogeneity has important implications, so the development of simple procedures for assessing the appropriate degree of heterogeneity within an aggregate model is important. Second, to our knowledge, this is the first aggregate analysis to address how individuals transition within the hours worked distribution. Previous analyses have instead only focused on how individuals transition between the states of employment and non-employment.

Having developed an empirically reasonable model that allows for adjustment along the intensive and extensive margin, we study the response of the model to business cycle shocks. For ease of exposition and transparency we focus on a shock that is commonly studied and hence consider aggregate shocks to productivity as the driving force behind business cycles. Of particular interest is to compare the model with an operative intensive margin to a model that exogenously abstracts from such adjustment. For this purpose we will consider the extensive-margin-only model of Chang and Kim (2007) as our benchmark model. Importantly, we assume that both models are calibrated to match the same aggregate targets and are subject to the same aggregate shocks. A key question is whether the model that exogenously shuts down the intensive margin is a good approximation to the behavior of the model that features an operative intensive margin. That is, we assess the extent to which abstracting from the intensive margin is a harmless simplification.

Intuition and previous work both suggest that the answer to this question depends on some of the underlying primitives of the economy, notably the willingness of individuals to substitute hours intertemporally and the extent of heterogeneity. For example, if individuals are not very willing to substitute labor intertemporally then intuitively there will be very little adjustment on the intensive margin and one might reasonably conjecture that it can be ignored. Because there remains some disagreement about the value of this curvature parameter we will consider a range of values consistent with the range of estimates found in the literature. Existing work also suggests that the extent of heterogeneity matters. For example, contrasting the results in Prescott et al. (2009) and Rogerson and Wallenius (2009) we see that when there is no heterogeneity all adjustment takes place along the extensive margin even though the intensive margin is available. For this reason we also consider specifications that feature different degrees of heterogeneity.

For each specification of curvature and heterogeneity, we calibrate the models to the same aggregate targets and then compare the business cycle fluctuations of our model with that of the benchmark model in which the intensive margin is exogenously shut down. Our main finding is that abstracting from the intensive margin can significantly distort inference regarding the volatility of aggregate hours. Moreover, the direction of the distortion depends on the underlying primitives. Surprisingly, even if variation along the intensive margin is very small, explicit modeling of the intensive margin can have a large impact on the volatility of aggregate hours. The presence of an intensive margin also dampens the effects of increased heterogeneity on aggregate hours volatility in an important way. We conclude that abstracting from the intensive margin is a serious issue for analyses that seek to understand the effect of heterogeneity on the magnitude of aggregate fluctuations.

Our analysis also offers an important message regarding the determination of intensive and extensive margin elasticities. While it is intuitive to think that the intensive margin elasticity is determined by the curvature parameter whereas the extensive margin elasticity is determined by the properties of heterogeneity, we find that intensive and extensive elasticities are not independent of each other and so must be considered jointly. That is, heterogeneity matters for the extent of adjustment along the intensive margin, and curvature matters for the extent of adjustment along the extensive margin. Our analysis also provides a mapping from the specification of the underlying primitives of heterogeneity and curvature in preferences over hours of work into the curvature parameter in a representative-agent model that would generate the same volatility in aggregate hours. Relative to a model that only features an extensive margin, our model produces a much smaller range of curvature values for the stand-in household, with values of Frisch elasticity in the range between 1.0 and 2.0.

Our paper is related to several in the literature. Relative to the business cycle analysis of Chang and Kim (2007), we add an intensive margin. Relative to the tax analysis of Rogerson and Wallenius (2009), we consider a richer environment in terms of heterogeneity, and allow for uncertainty and incomplete markets. The model that is probably closest to ours is Erosa et al. (2011). Like us, they build a model that features heterogeneity and incomplete markets and allows for labor supply adjustment along the intensive and extensive margin. They adopt a life cycle structure, consider a richer environment in terms of sources of heterogeneity, and calibrate the model to evaluate the labor supply elasticity of working age males in a partial equilibrium context. While their framework is richer, our somewhat more abstract model is more tractable and lends itself more readily to general equilibrium analyses. But most importantly, we investigate very different questions: whereas they focus on calibrating their model to evaluate its aggregate labor supply elasticity, we focus on the role that explicit modeling of the intensive margin plays in shaping the nature of aggregate labor market fluctuations. We therefore view these two pieces of work as complementary.

The paper is organized as follows: Section 2 specifies the model. Section 3 calibrates the different specifications of the model economy and Section 4 considers the steady state properties of the various specifications. Section 5 studies the business cycle properties of the model. Section 6 concludes.

2. Models

In this section we describe the two models that we will be comparing in our quantitative business cycle analysis. The benchmark model will be a model that only features adjustment along the extensive margin, and is identical to the model in Chang and Kim (2007). The other model that we consider extends Chang and Kim (2007) in the spirit of Prescott et al. (2009) to also allow for adjustment along the intensive margin. Because the benchmark model can be viewed as a special case of the more general model, we begin by describing the model with active adjustment along both the intensive and extensive margins.

2.1. Model with Intensive and Extensive Margin Adjustment

There is a unit measure of ex-ante identical infinitely lived individuals. Each individual has preferences over streams of consumption (c_t) and hours of work (h_t) given by:

$$\sum_{t=0}^{\infty} \beta^t \left\{ \log c_t - B \frac{h_t^{1+1/\gamma}}{1+1/\gamma} \right\}$$

where $0 < \beta < 1$, $B > 0$ and $\gamma > 0$. Here, γ represents a curvature parameter for the willingness to substitute hours over time.

There is an aggregate Cobb-Douglas production function that produces output using inputs of labor services (L_t) and capital services (K_t) and is subject to TFP shocks (λ_t):

$$Y_t = \lambda_t L_t^{\alpha} K_t^{1-\alpha}.$$

The aggregate productivity λ_t evolves with a transition probability distribution function $\pi_\lambda(\lambda'|\lambda) = \Pr(\lambda_{t+1} \leq \lambda'|\lambda_t = \lambda)$. In our quantitative analysis we will assume that λ_t follows an AR(1) process in logs:

$$\ln \lambda_{t+1} = \rho_\lambda \ln \lambda_t + \varepsilon_{\lambda t},\ \ \varepsilon_{\lambda t} \sim N(0, \sigma_\lambda^2).$$

Output can be used for either consumption or investment, and capital depreciates at rate δ each period.

Two features influence the mapping from time devoted to work to labor services. The first is that individuals are subject to idiosyncratic productivity shocks, denoted by x_t. The stochastic evolution of x_t is described by the transition probability distribution function $\pi_x(x'|x) = \Pr(x_{t+1} \leq x'|x_t = x)$. In our quantitative work we will also assume that x_t follows an AR(1) process in logs:

$$\ln x_{t+1} = \rho_x \ln x_t + \varepsilon_{xt},\ \varepsilon_{xt} \sim N(0, \sigma_x^2).$$

The second feature is a non-convexity associated with such factors as set-up costs, supervisory time and/or the need to coordinate with other workers. If an individual with idiosyncratic productivity x_t devotes h_t units of time to market work, this will generate $x_t g(h_t)$ units of labor services. Following Prescott et al. (2009) and Rogerson and Wallenius (2009), we assume that $g(.)$ takes the following simple form[1]:

$$g(h_t) = \max\left\{0, h_t - \bar{h}\right\},\ h_t \in [0,\, 1].$$

Following Bewley (1986), Huggett (1993), and Aiyagari (1994) we assume that markets are incomplete in the sense that there are no markets for insurance and the only asset is physical capital. Individuals trade claims to physical capital, and these claims are denoted by a. Additionally, there is an exogenous borrowing constraint that limits the amount of debt that an individual can acquire:

$$a_t \geq \bar{a}$$

In each period t there is a market for units of labor services, with price w_t, and a rental market for capital services, with price $r_t + \delta$, so that r_t is the rate of return to capital. When a worker of productivity x_t devotes h_t units of time to market work, the resulting labor earnings are $w_t x_t g\left(h_t\right)$.

Our model assumes that all (exogenous) heterogeneity occurs along one dimension–that of productivity in market work. More generally, one could also imagine that individuals differ along a second dimension, which could be thought of as the return to non-market work, either in the form of differences in the value of leisure time or in the productivity of non-market work. From the perspective of market labor supply choices, what really matters is the relative return to market work. While we could have introduced a second idiosyncratic shock, to maintain parsimony we have opted for a single idiosyncratic shock that we will interpret as a composite shock. It will be relevant to keep this in mind when we discuss calibration.

[1] French (2005) considers an alternative and smoother specification in which the wage per unit of time is smoothly increasing in the number of hours worked.

We formulate equilibrium recursively. The individual state variables are beginning of period assets (a) and current idiosyncratic productivity (x), and the aggregate state variables will be the current aggregate productivity shock (λ) and a measure μ over the individual state variables (a, x). Prices are functions of the aggregate state: $w(\lambda, \mu)$ and $r(\lambda, \mu)$, and the equilibrium law of motion for μ is given by $\mu' = T(\lambda, \mu)$.

The value function for a worker, denoted by V, is:

$$V\left(a, x\ ; \lambda, \mu\right) = \max_{c, a', h} \left\{ \log\left(c\right) - B\frac{h^{1+\frac{1}{\gamma}}}{1+\frac{1}{\gamma}} + \beta E\left[\ V\left(a', x'; \lambda', T\left(\lambda, \mu\right)\right) \mid x,\ \lambda\ \right] \right\}$$

subject to

$$c = w(\lambda, \mu)x \cdot \max\left\{0, h - \bar{h}\right\} + \left(1 + r(\lambda, \mu)\right) a - a'$$

$$c \geq 0,\ a' \geq \bar{a},\ 0 \leq h \leq 1$$

An equilibrium consists of a value function $V\left(a, x; \lambda, \mu\right)$, individual decision rules $c\left(a, x; \lambda, \mu\right)$, $a'\left(a, x; \lambda, \mu\right)$, $h\left(a, x; \lambda, \mu\right)$, aggregate inputs $\left\{K\left(\lambda, \mu\right),\ L\left(\lambda, \mu\right)\right\}$, factor prices $\left\{w\left(\lambda, \mu\right), r\left(\lambda, \mu\right)\right\}$, and a law of motion $T\left(\lambda, \mu\right)$ such that

1. Individuals optimize:

   Given factor prices, individual decision rules solve value function.

2. The representative firm maximizes profits : For all (λ, μ)

$$\begin{aligned} w\left(\lambda, \mu\right) &= F_1\left(L\left(\lambda, \mu\right),\ K\left(\lambda, \mu\right),\ \lambda\right) \\ r\left(\lambda, \mu\right) &= F_2\left(L\left(\lambda, \mu\right),\ K\left(\lambda, \mu\right),\ \lambda\right) - \delta \end{aligned}$$

3. The goods market clears : For all (λ, μ)

$$\int \left\{a' + c\right\} d\mu = Y + \left(1 - \delta\right) K$$

4. Factor markets clear :

$$\begin{aligned} L\left(\lambda, \mu\right) &= \int xg\left(\ h\left(a, x; \lambda, \mu\right)\ \right) d\mu \\ K\left(\lambda, \mu\right) &= \int a d\mu \end{aligned}$$

5. Individual and aggregate behaviors are consistent :

$$\mu'\left(A^0, X^0\right) = \int_{A^0, X^0} \left\{ \int_{\mathcal{A}, \mathcal{X}} 1\left[a' = a'\left(a, x; \lambda, \mu\right)\right] d\pi_x\left(x'|x\right) d\mu \right\} da' dx'$$

   for all $A^0 \subset \mathcal{A},\ X^0 \subset \mathcal{X}$.

2.2. Benchmark Model With Extensive Margin Adjustment Only

The model of Chang and Kim (2007) can be viewed as a special case of the above model in which we set the fixed cost $\bar{h} = 0$ and restrict the choice set for hours to a set with only two elements, 0 and $\hat{h}$. It is important to note that the parameter $\bar{h}$ acts like a fixed cost associated with working. It is well known (see, e.g., Cogan (1981)) that some kind of fixed cost is critical to generate the type of hours worked distributions that we see in reality. In a model that assumes exogenously that individuals can only choose one value for hours of market work the need for this fixed cost is obviated and so is not included.

When the choice of hours is restricted to be either 0 or $\hat{h}$, the parameter γ plays no separate role from the parameter B, since only the value $B\frac{\gamma}{1+\gamma}\hat{h}^{\frac{1+\gamma}{\gamma}}$ matters. Without loss of generality we will normalize the value of γ to be equal to one for the benchmark model.

3. Calibration

Our objective is to understand the extent to which the benchmark model with only an extensive margin does a reasonable job of capturing the properties of business cycle fluctuations in labor market variables that occur in the more general model that allows for intensive margin adjustment. That is, we want to ask whether a modeler might reasonably choose to work with the (simpler) benchmark model even though he or she recognizes that there is indeed some adjustment along the intensive margin in the data. In principle one would expect that the answer to this question may well depend on the values of some key parameters. Two special cases suggest two obvious features of the economy that will likely be important. As one extreme case, Prescott et al. (2009) show in their model that if all workers have identical productivity at each point in time, all adjustment will occur along the extensive margin even though the model explicitly allows for adjustment along the intensive margin. In this case there would be no loss in generality in considering the model that restricts adjustment to the extensive margin. Rogerson and Wallenius (2009) show that this result does not hold if there is heterogeneity in productivity. The second extreme case corresponds to $\gamma = 0$. In this case all individuals who work will choose the same hours of work, and again the model will be observationally equivalent to a model that only has an extensive margin. But as long as γ is positive, variation in individual productivity will lead to variation in hours worked.

Motivated by the above discussion, we will consider economies that differ in terms of the extent of heterogeneity and in the value of γ. One simple way to vary the amount of heterogeneity in the economy is to vary the standard deviation of the innovations to the idiosyncratic shock process, σ_x, and this is the approach that we follow. If there were definitive estimates for σ_x and γ it would perhaps be sufficient to just carry out our exercise for these specific values. However, we do not think that this is the case, and as a result consider a range for each of these parameters. More generally, examining a range of values is still of interest to get a clearer understanding of the underlying economics.

We noted earlier that our idiosyncratic shock is best thought of a composite shock that reflects the net effect of idiosyncratic shocks on the relative return to working in the market versus not working. A reasonable lower bound on the size of these shocks is provided by the literature that estimates idiosyncratic shocks to wages.[2] A sizable literature has done this for prime age males, including, for example, Card (1994), Floden and Linde (2001), French (2005), Chang and Kim (2006), and Heathcote et al. (2008). While there is some variation across studies, the consensus is that these shocks are large and persistent. Guided by these empirical studies, for one of our specifications we set $\rho_x = 0.975$ and $\sigma_x = 0.165$ at quarterly frequency.[3]

As noted above, we view this as a reasonable lower bound on the extent of heterogeneity induced by idiosyncratic shocks. As an upper bound we consider a specification with the same persistence, i.e., $\rho_x = .975$, but double the standard deviation of the innovations so that $\sigma_x = .330$. As we document later, the reason we view this is a reasonable upper bound is that it generates a dispersion in hours worked in steady state that exceeds what is found in the data. We also consider one intermediate value and so assume that σ_x takes on values in the set $\{0.165, 0.2475, 0.330\}$.

Motivation for the set of values considered for γ comes from large literature that has sought to estimate this parameter. Chetty (2012) argues that an empirically reasonable value for this elasticity parameter is in the range of $.40 - .50$. Early estimates such as MaCurdy (1981) suggested values around .25, and we include this value as a conservative lower estimate. Using a very different method, Pistaferri (2006) found an estimate around .75. Rogerson and Wallenius (2013) provide evidence for values of γ that are 1 or larger, and as discussed in Keane and Rogerson (2012), there are additional factors that Chetty abstracts from that would suggest higher values. For this reason we think values of γ as high as 1.00 are still quite plausible, and consider 1.50 to be a very generous upper bound. We therefore assume that γ takes on values in the set $\{0.25, 0.5, 1.0, 1.5\}$.

[2]This is reasonable as long as other idiosyncratic shocks are not perfectly negatively correlated with idiosyncratic wage shocks.

[3]Note that all of the papers previously cited estimated shocks based on annual data, so that our benchmark values need to be converted to annualized values when comparing them to the literature. The values just reported correspond to the estimates in Floden and Linde (2001).

Given values for σ_x and γ, we now describe how we calibrate the general model. As is standard in the business cycle literature we assume that each period corresponds to one quarter. Many of our parameters are standard in the literature and so we set them to be in line with previous studies. In particular the Cobb-Douglas technology parameter α is set to 0.64 and the depreciation rate, δ, is set to .025, effectively targeting labor's share of income and the investment to output ratio. For the aggregate technology shocks we set $\rho_\lambda = 0.95$, and $\sigma_\lambda = 0.007$. These values are all independent of γ and σ_x.

There are four additional parameters to calibrate: β, B, $\bar{a}$, and $\bar{h}$. Given values for γ and σ_x, we choose these values to match four moments in the steady state equilibrium: the (quarterly) rate of return to capital is 1%, the employment rate is 70%, the borrowing limit is equal to two quarters' earnings of a worker with productivity equal to the mean of all employed workers, and average hours (conditional on working) is 1/3.[4] Note that we are following the standard practice in the business cycle literature of recalibrating the model as we vary the values of γ and σ_x so that each of the 12 specifications match the same aggregate targets.

We want to consider how inference would be affected if the researcher were to start with the benchmark model that abstracts from choice along the intensive margin rather than the more general model. In doing so we want the researcher to continue to match the same targets. We do this by fixing h to be 1/3 for all workers, and removing the parameter $\bar{h}$ from the calibration exercise. All other aspects of the calibration exercise remain the same. In particular, the values of B, β, and $\bar{a}$ are again chosen so that in the steady state the (quarterly) rate of return to capital is 1%, the employment rate is 70%, and the borrowing limit is equal to two quarters' earnings of a worker with productivity equal to the mean of all employed workers. The parameter $\bar{h}$ does not exist in this specification, but we have chosen the exogenous fixed working time so as to hit the target of $h = 1/3$ conditional on working.[5] As noted earlier, the value of γ is not relevant for this economy, so there are effectively three specifications for the benchmark economy which differ in terms of σ_x only.

Table 1 summarizes the parameter values that are held constant across specifications as well as

[4]With a quarterly employment rate of 70%, the average annual employment rate in our model (i.e., fraction of individuals who work at least one quarter during a year) is 76.7%. This corresponds to the average annual employment rate in the PSID over the period 1968-1998 for household heads and spouses with ages between 18 and 65. We use a cutoff of 240 annual hours as the threshold for employment, i.e., we treat individuals with less than 240 annual hours worked as not employed.

[5]There is one exception to note. When $\sigma_x = .330$ the natural borrowing constraint for the economy with only an extensive margin is tighter than the borrowing constraint used previously, so in this case the borrowing limit is set to zero. The borrowing constraint is tighter in an economy without an intensive margin since it precludes people in low productivity states with high debt from working more hours as a way to generate more income in these states.

those that vary across specifications.

4. Properties of Steady State

In this section we consider some of the properties of the steady state equilibrium. There are two main objectives of this section. The first is to demonstrate that although our model is highly stylized, it is able to capture many features of the heterogeneity in wealth, earnings and hours worked found in the data. The second objective is to examine how different values for γ and σ_x influence the model's ability to account for these features in the data.

4.1. The Hours Worked Distribution

As noted in the calibration section, all of our model specifications are calibrated so as to generate the same fraction of people employed and the same average hours for workers conditional on employment. In this subsection we examine the extent to which the model can account for the distribution of hours worked among workers, and how this distribution varies across the various model specifications. We begin with Table 2, which reports standard deviations of the steady state distribution of annual hours of work, conditional on working. We compute this measure at the annual level because it is not available at the quarterly level in the PSID. Our sample is all household heads and spouses between the ages of 18 and 65 during the period 1968-1997. In the data there are some individuals who work very few hours during the year. We therefore classify a worker as employed if he or she works at least 240 hours per year, and treat those with less than 240 annual hours as having zero hours.[6] In the model, an individual is classified as employed if he or she has positive hours for at least one quarter during the year. While for some of the subsequent analysis we will utilize the panel nature of the PSID, this feature is not essential to this calculation, and so as a robustness check we also include a measure based on the CPS.[7] We normalize annual hours in the CPS and PSID so that average annual hours is the same as in the economy with $\gamma = 1$ and $\sigma_x = .165$.[8]

The standard deviations in the PSID and CPS are fairly comparable—.418 in the former and

[6] This adjustment affects relatively few individuals and our results are not very sensitive to variation in this cutoff.

[7] Annual Hours for the CPS data is obtained by multiplying "Average Hours per Week" and "Number of Weeks Worked".

[8] Average annual hours do not vary much across model specifications, so we do not renormalize for comparison with each specification.

.454 in the latter. Acknowledging that there is likely to be some measurement error in hours, the true dispersion in hours will be less than indicated by these values. The standard deviations for the 12 model specifications range from .301 to .480. Consistent with intuition the hours dispersion is increasing in γ and in σ_x. When $\sigma_x = .165$, the model cannot generate sufficient dispersion in hours even with a relatively large value of γ. However, when $\sigma_x = .330$ the model is able to match the dispersion found in the PSID as long as γ is around 1.0 or larger. As noted earlier, this motivates our choice of $\sigma_x = .330$ as a reasonable upper bound on the extent of heterogeneity in our model.[9]

To examine the implications for the distribution of hours worked in somewhat more detail, we next look at the average hours worked at various percentiles of the hours distribution. Figure 1 plots these values for two of the model specifications (the high and low values of σ_x with $\gamma = 1$) as well as the corresponding values for the PSID.

Consistent with the message based on looking simply at the standard deviations, we see that the specification with $\gamma = 1$ and $\sigma_x = .330$ does a much better job of tracking the empirical hours distributions than does the specification with the lower value of σ_x. In fact, this specification tracks the distribution in the PSID quite well.

4.2. Employment and Hours Transitions

Having assessed the model's ability to account for the distribution of individuals between employment and non-employment as well as the distribution of hours among employed workers, we next examine the model's ability to account for the movement of individuals within the hours worked distribution, including transitions into and out of employment.

Our data on transitions comes from the PSID and so are based on annual measures. We begin by looking at transitions into and out of employment. Table 3 shows the distribution of individuals across different combinations of employment states in consecutive years for the PSID and one of our model specifications ($\gamma = 1$ and $\sigma_x = .165$). We only report statistics for one of the model specifications because it turns out that these statistics are virtually identical across all 12 specifications.

The model does a good job of accounting for transitions into and out of employment, though employment status is somewhat less persistent in the model than in the data, as evidenced by the fact that the model has a greater share of workers changing employment status across consecutive

[9] Table 2 might lead one to consider even higher values for σ_x combined with smaller values for γ. However, as we show in a later subsection, this leads to an estimated Frisch elasticity that is well below .40, and so we do not consider this region of parameter space.

years than does the PSID. However, this difference is relatively small. For example, in the PSID, the persistence of employment (i.e., the probability of being employed next year conditional on being employed this year) is .74 whereas this probability is .73 in the model.

Next we consider the transition matrices for annual hours worked between years t and $t+1$. Table 4 reports the transition rates between quintiles of the hours worked distribution (and non-employment) from the PSID and the two model specifications in which $\gamma = 1$ and the two extreme values of σ_x. The transition matrices for different values of γ holding σ_x constant turn out to be quite similar, so in the interest of space we do not include them. As we will see, although changes in σ_x do produce some quantitative differences, most notably in the degree of persistence, the basic patterns are also not affected much by changes in σ_x.

We start by noting three prominent features in the data. First, annual hours of work exhibit a significant degree of persistence, especially for those workers with high hours of work. All of the diagonal elements for workers with positive hours are greater than 40%, and for the highest quintile this value exceeds 60%. Second, conditional on working in both periods and switching quintiles, the transition probabilities are monotone decreasing in the distance of the destination quintile from the originating quintile. Third, individuals who adjust along the extensive margin between years t and $t+1$ are disproportionately from the lowest quintile of the hours distribution. Specifically, for those individuals who work in year t but not in $t+1$, roughly two-thirds of them have annual hours of work in the lowest quintile in year t. Similarly, for those individuals who did not work in year t but did work in year $t+1$, roughly three-quarters of them have annual hours in the lowest quintile in year $t+1$.

Next we consider how the model fares in terms of reproducing these features. As Table 4 shows, both model specifications generate considerable persistence, though less in the economy with $\sigma_x = .33$. The average of the diagonal elements for those with positive hours in both periods is 49.48 in the data, 46.13 in the economy with $\sigma_x = .165$ and 38.48 in the economy with $\sigma_x = .33$. While both model specifications come close to matching the persistence in the highest quintile, the lowest quintile displays quite a bit more persistence in the data than in either model–49.45 versus 32.33 and 32.66. One possible explanation for this is the existence of a group of workers in the data who desire part-time work on a more permanent basis than captured by the idiosyncratic shocks in our model.

For the most part the model also matches the second observation noted above. Specifically, for the specification with $\sigma_x = .33$ the model has the same monotonicity property found in the data, whereas for the other specification there are a couple of values which violate the pattern.

Finally, the model does a good job of matching the nature of adjustment along the extensive margin. For both of the model specifications shown in the table, roughly three quarters of all transitions along the extensive margin involve workers who are in the lowest quintile of the hours distribution.

Overall, we take the fact that the model can match the salient features of transitions within the hours distributions to suggest that our parsimonious representation of individual heterogeneity is empirically reasonable.

4.3. Distributions of Wealth and Earnings

Chang and Kim (2007) and An et al. (2009) have previously shown that a model with idiosyncratic productivity shocks calibrated to micro data, incomplete markets and indivisible labor captures many quantitative features of the wealth and earnings distribution. It turns out that adding an intensive margin to the analysis has little impact along this dimension. As a result, when $\sigma_x = .165$ our wealth and earnings distributions look very similar to those in An et al. (2009). When we consider the specification with a much greater degree of heterogeneity, $\sigma_x = .330$, we get similar patterns qualitatively, but the model generates too much dispersion in earnings. In this section we document these properties.

Given that we calibrate our model using employment data from the PSID we think it is preferable to compare our model to data that is also based on the PSID. For this reason our primary source of information on the cross-sectional wealth and earnings measures are based on the 1984 PSID. As a robustness check we also report comparable figures for properties of the wealth distribution from the work of Diaz-Gimenez et al. (1997) that is based on the Survey of Consumer Finances (SCF). For the measures that we focus on the two data sets provide very similar answers, so this does not seem to be a major issue.[10]

Table 5 reports the Gini coefficients for both the wealth and earnings distributions in the eight of the different model specifications that we consider, as well as their corresponding values in the PSID and SCF.

A few patterns are evident. First, given a value for σ_x, the Gini coefficients for both the wealth and earnings distributions are (weakly) increasing in the value of γ. (Note that the extensive-margin-only case can be thought of as the limiting case as γ goes to zero.) This effect is intuitive;

[10]This is not the case if one focuses on the extreme upper tail of the wealth distribution, as it is well known that the PSID under-samples the upper extremes of the wealth distribution. However, as noted before, given our emphasis on labor supply, this extreme upper tail is not of primary concern.

a higher value of γ leads to greater intertemporal substitution of labor supply, so that individuals work more when productivity is high and less when productivity is low, thereby amplifying the direct effect of productivity on earnings. Given that individuals accumulate assets to smooth consumption in the face of fluctuations in earnings, greater fluctuations in earnings leads to greater dispersion in assets. Although the qualitative effects are intuitive, the main message from Table 3 is that the quantitative importance of these effects are quite small. While moving from $\gamma = 1.0$ to the extensive-margin-only case does generate modest changes in both Gini coefficients, the effect of changes in γ inside the interval $[.25, 1.5]$ is of second order importance for each measure.

It is also intuitive that higher values of σ_x would similarly lead to increases in the Gini coefficients for both wealth and earnings distributions. However, in contrast to the previous case, changes in σ_x for a given value of γ do generate first order effects on both measures, with the effect on the earnings Gini being almost twice as large as the effect on the wealth Gini.

Comparing the values in the various model specifications with their counterparts in the data, all of the model specifications seem to capture much of dispersion in the wealth and earnings distributions. If anything, the models generate too much dispersion in earnings, especially in comparison to what is found in the PSID. To look a bit deeper into the nature of the wealth and earnings distributions, Table 6 shows the wealth and earnings shares by quintiles of the wealth distribution. Because variation in γ turns out to be not very important quantitatively in terms of these distributions, we only report results for the two model specifications with the extreme values of σ_x and with $\gamma = 1$.

The basic message from Table 6 is that in addition to doing a reasonable job of accounting for the absolute amount of dispersion in wealth and earnings as captured by the Gini coefficient, the model also does a good job of accounting for the shape of these distributions. The specification with $\sigma_x = .165$ does a very good job of capturing the earnings shares, whereas consistent with Table 2, the specification with $\sigma_x = .330$ generates excessive concentration of earnings in the upper quintile of the wealth distribution. However, the specification with the higher value of σ_x is better able to capture the amount of wealth concentrated in the upper quintile. Analyzing the wealth shares by quintiles of the wealth distribution hides the extreme concentration of wealth at the very top of the distribution. It is well known (see, for example, Diaz-Gimenez et al. (1997)) that the model is not able to account for the concentration found in say the upper 1% of the wealth distribution. However, from the perspective of labor supply, accounting for the likes of Bill Gates is probably not of first order importance, and so we do not focus on the extreme upper part of the wealth distribution.

In summary, this subsection shows that all of the model specifications generate significant dispersion in earnings and wealth relative to the data. If anything, some of the model specifications generate too much dispersion. The nature of the dispersion is also empirically reasonable, in terms of matching earnings and wealth shares by quintiles of the wealth distribution. We conclude that adding an intensive margin of labor supply to the previous analyses of Chang and Kim (2006, 2007) does not have first order effects on the earnings and wealth distributions.

5. Business Cycle Analysis

In this section we study the business cycle properties of our model economies. We solve the equilibrium of the model using the method proposed by Krusell and Smith (1998). Details are included in the appendix. In this section we focus on aggregate statistics, and since aggregate data is available at quarterly frequency, we also compute model statistics using quarterly data. As is standard, we take logs and then HP filter (with the smoothing parameter of 1,600) the simulated (as well as the actual data) series before computing statistics.

Table 7 reports the properties of output and labor market variables from our models. Aggregate hours worked, employment (extensive margin), hours per worker (intensive margin) and aggregate efficiency units of labor services are denoted as H, E, h, and L respectively. By definition, $H = E \times h$. Because the behavior of consumption and investment is basically the same as in standard real-business-cycle exercises we do not report statistics for these variables.

Next we turn to the properties of fluctuations in the models. We emphasize that in all of these model economies the driving force behind aggregate fluctuations is identical, i.e., the parameters of the technology shock process are held constant across all specifications. Hence, to the extent that fluctuations are different in the various model economies, it is the result of how the models lead to different propagation of these shocks. While our focus is on labor market variables, we first note that consistent with many previous exercises, the technology shocks that we feed into our model generate output fluctuations that are between two-thirds and three-quarters of observed fluctuations in output. We will say more about the nature of these differences below, when we examine the nature of labor market fluctuations in more detail.

It is useful to first summarize the key patterns for the benchmark model that features only an extensive margin. There are three specifications of this model, differing in the extent of heterogeneity. The main pattern is that as σ_x increases from .165 to .33 the volatility of employment and aggregate hours falls very dramatically, from .93 to .30. The intuition for this result is simple and well-known. In steady state the optimal decision rules define a curve that divides the individual state space into two distinct regions, one with working as the optimal decision and the other with not working as the optimal decision. A positive aggregate productivity shock increases the wage rate and encourages workers who were previously slightly below the work threshold to engage in market work. The size of the employment response depends on the density of individuals near the boundary, and this in turn is influenced by the extent of heterogeneity. A larger value of σ_x tends to spread out the distribution of workers across the state space, thereby lowering this density and decreasing the employment response. The flip side of this is that the wage per efficiency unit of labor will fluctuate more.

While the qualitative pattern just noted is intuitive, we want to emphasize the magnitude of the effects. Moving from $\sigma_x = .165$ to $\sigma_x = .33$ reduces volatility in aggregate hours by a factor of three. Whereas the economy with $\sigma_x = .165$ accounts for roughly two thirds of the volatility of aggregate hours relative to output, the $\sigma_x = .33$ specification accounts for less than a third of this measure. Finally, note also that in all three specifications of the benchmark model, the volatility of efficiency units of labor is less than the volatility in aggregate hours (and employment). The reason for this is that the fluctuations in employment are dominated by individuals with idiosyncratic productivity that is lower than the mean idiosyncratic productivity of employed individuals. While a positive aggregate productivity shock can induce some high wealth-high idiosyncratic productivity individuals to begin working, this group does not dominate the movements into and out of employment. As σ_x increases and the extent of heterogeneity increases, the gap between the volatility in hours and efficiency units becomes larger.

We now ask to what extent this benchmark model accurately captures labor market fluctuations in the more general models that include an operative intensive margin, i.e., to what extent does it matter that one abstracts from adjustment along the intensive margin. Before turning to specific results, it is instructive to return to the extreme special cases that we mentioned earlier. Specifically, we argued that either little heterogeneity or a very low value of γ might suggest that it is very innocuous to abstract from adjustment along the intensive margin. With this in mind, consider the general model with $\sigma_x = .165$ and $\gamma = .25$. In the calibration section we argued that both of these were conservative lower bounds and hence are a relevant case to consider. Table 7 reveals a striking result: whereas the benchmark model generates a volatility of aggregate hours equal to .93 when $\sigma_x = .165$, the more general model with $\sigma_x = .165$ and $\gamma = .25$ only generates a volatility of aggregate hours equal to .68, a reduction of more than twenty-five percent.

To better understand this result it is instructive to first look at the corresponding fluctuations in efficiency units. Interestingly, the volatility of efficiency units of labor is virtually the same in both cases. So why is the difference between fluctuations in hours and efficiency units so pronounced between these two economies? This is intimately related to the two features that one needs to generate adjustment along both margins, namely heterogeneity in conjunction with fixed costs. Specifically, as already noted, the individuals who move between employment and non-employment in response to aggregate shocks tend to be lower productivity. In the extensive-margin-only case, these individuals will work the same number of hours as existing workers, but in the model with an active intensive margin, these individuals will work fewer hours than the individuals who do not move in and out over the cycle. The presence of fixed costs implies that the efficiency units per hour of work are lower at lower hours of work, so that workers moving into and out of employment over the business cycle contribute much smaller changes in efficiency units than do existing workers. It follows that a given change in the wage per efficiency unit will elicit a smaller response along the extensive margin. Note that the wage per efficiency unit fluctuates very similarly in the two economies but that employment fluctuates much less in the more general model. Moreover, when existing workers increase their hours along the intensive margin, the marginal change in efficiency units is much greater. As a result, even a small increase along the intensive margin for existing workers is much more powerful in changing efficiency units. The final outcome is that although efficiency units fluctuate virtually the same as in the benchmark model, this is achieved in a very different fashion, thereby manifesting itself as very different fluctuations in aggregate hours.

It is important to emphasize that there is a striking difference between the volatility of hours in the two economies despite the fact that in the more general economy, the volatility of hours per worker is only .03 and hence is very close to zero. The key message here is that even if there is little volatility along the intensive margin, abstracting from it has important implications precisely because a model that has an active intensive margin will necessarily contain features that are important for the nature of labor market fluctuations. In this regard it is also of interest to note that this first comparison also contradicts another simple piece of intuition. Specifically, when one adds an intensive margin choice to a model that only has an extensive margin choice one might conjecture that adding additional margins of adjustment could only lead to greater adjustment in hours. But this intuition neglects the fact that generating an active intensive margin requires additional changes in the model, and these changes affect the nature of labor market adjustments.

Holding the value of σ_x constant at .165 and increasing the value of γ delivers fairly intuitive results: higher values of γ make adjustment along the intensive margin less costly, and as a result we see that not only does the relative volatility of hours per worker increase relative to employment, but also that there is a small increase in the volatility of aggregate efficiency units. Nonetheless, the volatility of aggregate hours is effectively unchanged as we vary γ at this value of σ_x. In particular, when $\sigma_x = .165$, the model with an operative intensive margin implies a much smaller volatility of aggregate hours than the extensive margin only model even if we assume a very large value for γ such as $\gamma = 1.5$.

Next we focus on the case in which $\sigma_x = .330$. As we noted before, in the benchmark model there is a dramatic decrease in the volatility of aggregate hours when we move from $\sigma_x = .165$ to $\sigma_x = .33$. Whereas we found that allowing for an active intensive margin significantly decreased the extent of hours volatility when $\sigma_x = .165$, we now see the opposite effect. For example, for γ in the range of .50 to 1.00, the volatility of aggregate hours is more than fifty percent higher for the model that allows for an operative intensive margin. Even for the conservative case of $\gamma = .25$ we see that the volatility of aggregate hours is almost one third larger relative to the case with only adjustment along the extensive margin.

It follows that the effect of increased heterogeneity on the volatility of aggregate hours is muted considerably by allowing for an intensive margin. For example, whereas moving from $\sigma_x = .165$ to $\sigma_x = .33$ in the benchmark model lowers the standard deviation of aggregate hours by .63, the analogous drop when $\gamma = .5$ is only .20. There are two factors at work here: first, the drop in the volatility in employment is lower, .31 versus .63, and second, there is a partially offsetting increase in the volatility of hours per worker, equal to .11. Loosely speaking, greater heterogeneity dampens adjustment on the extensive margin, and as a result we see more adjustment along the other margin. If this margin is shut down we overstate the loss in volatility.

Our analysis also offers an important message concerning the determinants of fluctuations along the intensive and extensive margins. In particular, we start from the notions, implicit in much of the recent literature, that fluctuations along the intensive margin are determined mostly (if not exclusively) by the value of the preference parameter γ, whereas fluctuations along the extensive margin are determined mostly (if not exclusively) by the extent of heterogeneity, which in our model is reflected in the value of σ_x. Intuitively, fluctuations along the intensive margin are increasing in the value of γ, while fluctuations along the extensive margin are decreasing in the value of σ_x.

While the results in Table 7 provide partial support for both of these notions, they also reveal that these notions reflect an oversimplification that is potentially very misleading from a quantitative perspective. For example, consistent with the first notion above, we see that for a given value of σ_x, increases in γ lead to greater fluctuations along the intensive margin. However, the notion that γ is the dominant, let alone the exclusive factor that determines this response is strongly contradicted by the results in Table 7. Specifically, starting from the specification in which $\gamma = 0.5$ and $\sigma_x = .165$, we see that increasing γ to 1.5 leads to almost a tripling of the response along the intensive margin. However, this same effect occurs if we keep the value of γ fixed at 0.50 but we instead increase the value of σ_x to .330. To the best of our knowledge, we are the first to point out that the aggregate fluctuations along the intensive margin are affected in a quantitatively critical way by the amount of heterogeneity in the economy.

Similarly, starting from the specification $\gamma = .5$ and $\sigma_x = .165$, we note that moving to the specification with $\sigma_x = .33$ leads to roughly a 50% reduction in fluctuations along the extensive margin. However, if we instead kept the value of σ_x unchanged at .165 and instead increased the value of γ to 1.5, we would still have a decrease in fluctuations along the extensive margin of almost 20%. It follows that one should not think of "intensive margin elasticities" and "extensive margin elasticities" as two independent parameters.

Our analysis also allows us to construct a mapping from the values of γ and σ_x into the implied curvature parameter for a representative household that would generate the same fluctuations in hours worked. In particular, we consider a representative household model in which the household has period utility function given by:

$$\log c + \frac{B}{1+\frac{1}{\hat{\gamma}}} h^{1+\frac{1}{\hat{\gamma}}}$$

We calibrate the value of B so that hours worked in steady state are 1/3 and consider the same aggregate shock process as before. We then ask what value for $\hat{\gamma}$ provides the same volatility in aggregate hours as the model with heterogeneity and incomplete markets. Table 8 shows the results for both the general model and the model with only an extensive margin. The first point to note is that the range of implied values for $\hat{\gamma}$ is much smaller once one includes an intensive margin. Whereas the range is from .64 to 4.94 in the benchmark model with only an extensive margin, the range is basically from 1.0 to 2.0 for the models that include an intensive margin, even though we consider a wide range of values for the individual curvature parameter. Moreover, if we were to restrict attention to $\sigma_x = .33$ on the grounds that it generates an empirically reasonable amount of heterogeneity in hours, and also consider values of γ that lie between .5 and 1.0, we see that the implied representative agent elasticity parameter is quite tightly pinned down and is a bit larger than 1.0. Even with γ as large as 1.5 we see that the representative agent elasticity is still only as high as 1.45

6. Conclusion

Recent advances in modeling aggregate labor supply have emphasized the importance of accounting for adjustment along the intensive and extensive margins. Adjustment along the extensive margin has also been shown to depend on the extent of heterogeneity. In this paper we build a model in which individual labor supply adjusts along both the intensive and extensive margins in an environment that features heterogeneity and incomplete markets. We believe that this is the appropriate benchmark model for understanding the joint determination of adjustment along the two margins. We consider a family of specifications of this model that differ along two key dimensions: the value of the preference parameter that influences curvature of utility in hours of work, and the standard deviation of innovations in the idiosyncratic shock process, which in turn influences the extent of heterogeneity in the invariant distribution for idiosyncratic shocks.

We consider the ability of the various specifications of the model to account for key features of employment and hours worked in the cross section. We then use this model to consider labor supply responses to aggregate technology shocks. Three key findings emerge. First, in terms of fluctuations in aggregate hours, we find that abstracting from the intensive margin can be very misleading about the effect of heterogeneity. Second, extensive and intensive margin elasticities are jointly determined by both the preference parameter and the extent of heterogeneity. That is, one cannot speak of intensive and extensive margin elasticities as independent parameters of the economic environment. Third, using data on the cross-section distribution of hours worked, we find that our models generate volatility in aggregate hours that would correspond to those generated in a representative agent model where the household has a Frisch elasticity slightly larger than unity. More generally, the model with both intensive and extensive margins of adjustment produces much less sensitivity in aggregate volatility in response to changes in the underlying primitives.

References

[1] Aiyagari, S. Rao (1994): "Uninsured Idiosyncratic Risk and Aggregate Saving," *Quarterly Journal of Economics*, 109(3), 659-684.

[2] Altonji, Joseph G. (1986) "Intertemporal Substitution in Labor Supply : Evidence from Micro Data.", *Journal of Political Economy*, 94(3), s176-215.

[3] An, Sungbae, Yongsung Chang, and Sun-Bin Kim (2009): "Can a Representative Agent Model Represent a Heterogeneous Agent Economy?" *American Economic Journal - Macroeconomics*, 1, 29-54.

[4] Bewley, Truman (1986) "Stationary Monetary Equilibrium with a Continuum of Independently Fluctuating Consumers," in *Contributions to Mathematical Economics in Honor of Gerard Debreu.*

[5] Card, David (1994) "Intertemporal Labor Supply: An Assessment," in Christopher Sims, editor, *Advances in Econometrics, Sixth World Congress*, New Yor: Cambridge University Press.

[6] Chang, Yongsung, and Sun-Bin Kim (2006): "From Individual to Aggregate Labor Supply: A Quantitative Analysis Based on a Heterogeneous Agent Macroeconomy," *International Economic Review*, 47(1), 1-27.

[7] Chang, Yongsung, and Sun-Bin Kim (2007): "Heterogeneity and Aggregation: Implications for Labor-Market Fluctuations," *American Economic Review*, 97(5), 1939-1956.

[8] Chetty, Raj (2012) "Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply," *Econometrica*, 80(3), 969–1018.

[9] Cogan, John F. (1981) "Fixed Costs and Labor Supply," *Econometrica*, 49(4), 945–963.

[10] Díaz-Giménez, Javier, Vincenzo Quadrini, and José-Víctor. Ríos-Rull, (1997) "Dimensions of Inequality: Facts on the U.S. Distributions of Earnings, Income, and Wealth," *Minneapolis Fed Quarterly Review* 21 , 3-21.

[11] Erosa, A., L. Fuster and G. Kambourov (2011) "Towards a Micro-Founded Theory of Aggregate Labor Supply," mimeo, University of Toronto.

[12] Floden, Martin, and Jesper Linde, (2001) "Idiosyncratic Risk in the United States and Sweden: Is There a Role for Government Insurance?," *Review of Economic Dynamics*, 4 (2), 406-437.

[13] French, Eric (2005) "The Effects of Health, Wealth, and Wages on Labor Supply and Retirement Behavior," *Review of Economic Studies*, 72(2), 395-427.

[14] Hansen, Gary D. (1985) "Indivisible Labor and the Business Cycle," *Journal of Monetary Economics*, 16(3), 309-327.

[15] Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante (2008) "Insurance and Opportunities: A Welfare Analysis of Labor Market Risk," *Journal of Monetary Economics*, 55, 501–525.

[16] Heckman, James J. (1984) "Comments on the Ashenfelter and Kydland Papers." *Carnegie-Rochester Conference Series on Public Policy,* 21, 209-224.

[17] Hugget, Mark (1993) "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies," *Journal of Economic Dynamics and Control*, 17, 953-69.

[18] Keane, Michael, and Richard Rogerson. (2012) "Micro and Macro Labor Supply Elasticities: A Reassessment of Conventional Wisdom," *Journal of Economic Literature*, 50(2), 464-76.

[19] Krusell, Per, and Anthony A. Smith (1998) "Income and Wealth Heterogeneity in the Macroeconomy", *Journal of Political Economy*, 89(6), 1059-1085.

[20] Kydland, Finn E. and Edward C. Prescott (1982), "Time to build and aggregate fluctuatations," *Econometrica*, 50, 1345-79.

[21] MaCurdy, Thomas E. (1981): "An Empirical Model of Labor Supply in a Life-Cycle Setting," *Journal of Political Economy*, 106(5), 867-96.

[22] Prescott, Edward, Richard Rogerson, and Johanna Wallenius (2009) "Lifetime aggregate labor supply with endogenous workweek length," *Review of Economic Dynamics*, 12, 23-36.

[23] Rogerson, Richard, and Johanna Wallenius (2009) "Micro and Macro Elasticities in a Life Cycle Model with Taxes," *Journal of Economic Theory* 144, 2277-2293.

[24] Rogerson, Richard, and Johanna Wallenius (2013) "Nonconvexities, Retirement, and the Elasticity of Labor Supply," *American Economic Review*, 103(4), 1445-1462.

Table 1: Calibrated Parameter Values

A. Values Held Constant Across Specifications

α	δ	ρ_λ	σ_λ	ρ_x
0.64	0.025	0.95	0.007	0.975

B. Values That Vary Across Specifications

		β	B	$\bar{h}$	$\bar{a}$
$\sigma_x = .165$	$\gamma = 1.5$	.9770	10.72	.1132	$-.6029$
	$\gamma = 1.0$	.9773	18.47	.1490	$-.5012$
	$\gamma = 0.5$	.9776	82.70	.2099	$-.3328$
	$\gamma = 0.25$	.9778	1236.8	.2590	$-.1993$
	Ext	.9792	18.05	–	$-.4741$
$\sigma_x = .2475$	$\gamma = 1.5$	.9698	9.32	.0945	-1.0899
	$\gamma = 1.0$	.9702	15.93	.1318	$-.9075$
	$\gamma = 0.5$	.9708	70.67	.1968	$-.6041$
	$\gamma = 0.25$	.9712	1051.7	.2508	$-.3619$
	Ext	.9734	14.75	–	$-.8230$
$\sigma_x = .330$	$\gamma = 1.5$	.9640	8.03	0.070	-2.2845
	$\gamma = 1.0$	.9645	13.50	.1074	-1.9170
	$\gamma = 0.5$	.9654	58.63	.1777	-1.2876
	$\gamma = 0.25$	.9659	863.9	.2383	$-.7737$
	Ext	.9678	11.49	–	0.0

Notes: 'Ext' denotes the model specification with the extensive margin only.

Table 2: Standard Deviation of Annual Hours Conditional on Working

Data		Models				
			γ=1.5	γ=1	γ=.5	$\gamma = .25$
PSID	CPS					
.418	.454	$\sigma_x = .165$	.337	.318	.301	.295
		$\sigma_x = .2475$	.379	.347	.314	.299
		$\sigma_x = .330$	.480	.413	.342	.310

Notes: Annual hours in the CPS and PSID are normalized so that average annual hours is the same as in the economy with $\gamma = 1$ and $\sigma_x = .165$.

Table 3: Distribution of Employment Transitions

E_{t-1}	E_t	Shares in PSID	Shares in Model
1	1	.740	.727
1	0	.046	.053
0	1	.032	.053
0	0	.182	.168

Notes: Employment status at time t (E_t), is denoted by 1 (working) or 0 (not working).

Table 4: Transition Probability of Hours (Annual)

			$t+1$					
PSID			Not Work	1st	2nd	3rd	4th	5th
		Not Work	83.57	12.25	1.69	0.91	0.99	0.60
		1st	21.08	49.45	14.91	6.15	5.29	3.12
		2nd	4.77	15.40	45.77	18.27	11.15	4.63
	t	3rd	2.81	6.75	19.77	46.24	17.88	6.54
		4th	2.26	5.14	10.63	19.91	42.42	19.64
		5th	1.81	3.41	4.69	6.80	19.77	63.52
$\sigma_x = 0.165$			Not Work	1st	2nd	3rd	4th	5th
		Not Work	76.12	18.62	3.64	0.28	0.92	0.43
		1st	25.33	32.33	23.05	3.30	9.33	6.66
		2nd	4.48	15.11	35.75	20.77	16.88	7.02
	t	3rd	0.42	2.24	13.86	62.16	17.95	3.37
		4th	1.92	11.79	14.05	12.04	39.22	20.98
		5th	1.54	12.28	8.15	1.56	15.29	61.17
$\sigma_x = 0.330$			Not Work	1st	2nd	3rd	4th	5th
		Not Work	76.96	18.32	3.85	0.65	0.17	0.04
		1st	25.30	32.66	25.75	10.43	3.89	1.97
		2nd	5.32	21.01	32.37	25.91	10.07	5.32
	t	3rd	1.44	9.97	19.25	30.65	26.11	12.58
		4th	0.54	5.24	9.49	19.89	40.82	24.03
		5th	0.26	4.98	7.49	12.42	18.93	55.91

Table 5: Gini Coefficients for Wealth and Earnings Distributions

Data	PSID	SCF	Models	γ=1.5	γ=1	γ=.5	$\gamma = .25$	Ext
Wealth	.76	.78	$\sigma_x = .165$	.74	.74	.73	.72	.70
Earnings	.53	.63		.64	.63	.62	.61	.58
			$\sigma_x = .330$	.81	.80	.80	.80	.78
				.84	.83	.82	.82	.78

Table 6: Wealth and Earnings Shares by Wealth Quintiles

		1^{st}	2^{nd}	3^{rd}	4^{th}	5^{th}
Wealth Share	PSID	-0.52	0.50	5.06	18.74	76.22
	SCF	-0.39	1.74	5.72	13.43	79.49
	$\sigma_x = .165$	-1.18	-0.48	5.76	20.55	75.35
	$\sigma_x = .330$	-1.06	0.10	3.34	12.94	84.68
Earnings Share	PSID	7.51	11.31	18.72	24.21	38.23
	SCF	7.05	14.50	16.48	20.76	41.21
	$\sigma_x = .165$	6.24	14.00	17.85	23.69	38.21
	$\sigma_x = .330$	1.06	3.87	8.33	17.96	68.78

Table 7: Business Cycle Statistics

		σ_Y	σ_H	$\frac{\sigma_H}{\sigma_Y}$	σ_E	σ_h	σ_L	σ_w
Data	ES	1.72	1.69	.98	1.41	.48	-	-
	HS	1.72	1.59	.92	1.14	.63	-	-
$\gamma = 1.5$	$\sigma_x = .330$	1.28	.56	.43	.26	.34	.61	.73
	$\sigma_x = .2475$	1.35	.60	.44	.35	.26	.71	.69
	$\sigma_x = .165$	1.46	.68	.47	.51	.17	.89	.64
$\gamma = 1.0$	$\sigma_x = .330$	1.24	.49	.40	.27	.24	.54	.75
	$\sigma_x = .2475$	1.31	.54	.41	.36	.19	.65	.71
	$\sigma_x = .165$	1.44	.66	.46	.54	.12	.87	.65
$\gamma = 0.5$	$\sigma_x = .330$	1.19	.47	.40	.30	.17	.46	.77
	$\sigma_x = .2475$	1.27	.47	.37	.36	.11	.57	.73
	$\sigma_x = .165$	1.43	.67	.47	.61	.06	.85	.66
$\gamma = 0.25$	$\sigma_x = .330$	1.16	.39	.34	.30	.09	.41	.78
	$\sigma_x = .2475$	1.24	.42	.34	.37	.06	.54	.74
	$\sigma_x = .165$	1.42	.68	.48	.65	.03	.84	.67
Ext	$\sigma_x = .330$	1.01	.30	.30	.30	0	.16	.85
	$\sigma_x = .2475$	1.13	.49	.43	.49	0	.38	.79
	$\sigma_x = .165$	1.39	.93	.67	.93	0	.85	.71

Note: Total hours (H) = Employment (E) $\times$ Hours per worker (h). "ES" and "HS" denotes the data based on the Establishment Survey and Household Survey, respectively. The aggregate efficiency unit of labor is denoted by L. The wage rate for the efficiency unit of labor is denoted by w.

Table 8: Implied γ for a stand-in household

	$\sigma_x = .33$	$\sigma_x = .2475$	$\sigma_x = .165$
$\gamma = 1.5$	1.45	1.66	2.16
$\gamma = 1.0$	1.17	1.37	2.02
$\gamma = .50$	1.10	1.10	2.09
$\gamma = .25$	.86	.95	2.16
Ext	.64	1.17	4.94

Figure 1: Distribution of Hours Worked

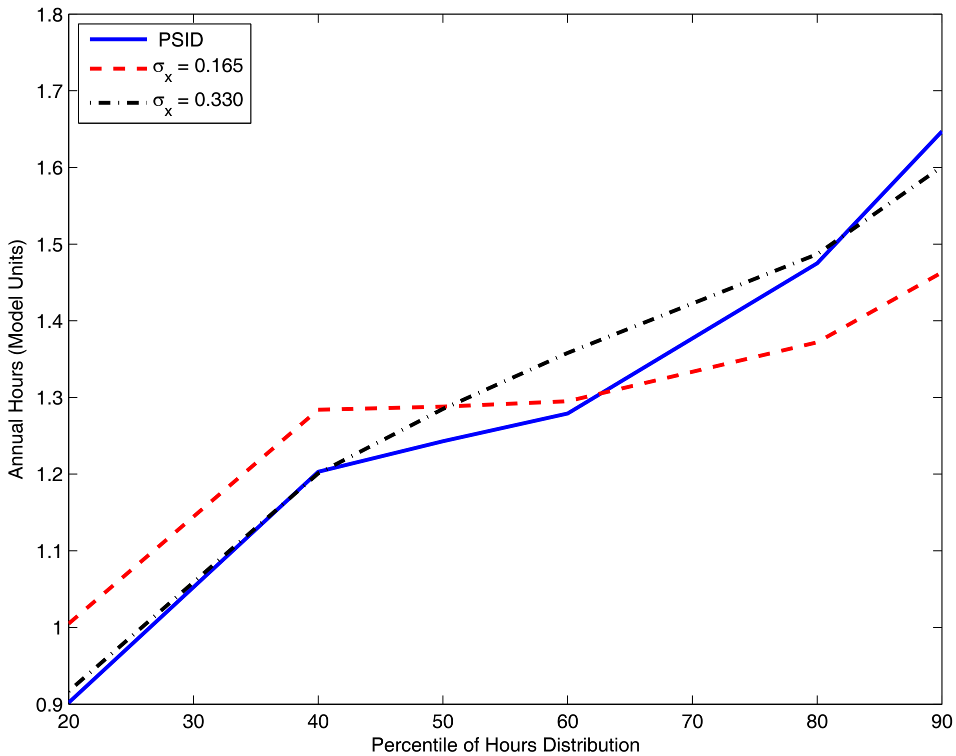

Notes: Annual hours in the PSID are normalized so that average annual hours is the same as in the economy with $\sigma_x = .165$.